

Introduction to EMC VNX2 Storage Efficiency Technologies

VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, & VNX8000

Abstract

This white paper provides an introduction of the VNX Storage Efficiency features on the EMC® VNX™ 2 storage systems. It provides an overview of the technology, and describes how the storage efficiency features are implemented on the VNX2 series to help you leverage daily storage activities.

March 2016

Copyright © 2016 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

VMware is a registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number H12197.4

Table of Contents

Executive summary.....	4
Audience.....	5
Introduction	5
Storage Efficiency Technologies overview.....	5
Unified Virtual Provisioning	6
FAST VP	7
Multicore FAST Cache	7
File Deduplication and Compression	7
Block Deduplication	8
VNX Block Deduplication for VNX File	8
Block Compression	8
Multi-dimensional storage efficiency.....	9
Capacity optimization.....	10
Performance optimization	14
Green Initiatives	16
Simplicity	18
Conclusion.....	20
References	20

Executive summary

Data is growing at a tremendous rate. In fewer than ten years, the amount of data managed by data centers is expected to grow by over 60 percent Compounded Annual Growth Rate (CAGR). Projections indicate that by 2020, data centers will need to store 14 times the amount of data stored in 2010.

Over the same period, the number of servers is expected to grow rapidly, mostly with virtual servers. In addition, processing power is expected to grow tenfold. The latest Intel multi-core architecture available in the EMC VNX2 systems is a perfect example of this rapid growth.

While hard drive density is also projected to increase, the rate of that increase is not fast enough to keep up with the amount of data being generated. At the same time, the available levels of IT budgets and staffing are projected to remain essentially flat over the same time period.

Figure 1 illustrates the Data Center dynamics from 2012 to 2020.




Figure 1. Data center dynamics 2012 to 2020¹

This rapid data growth means that, despite continued reductions in the cost per TB, total investments in storage will grow significantly over the next 10 years. Optimizing resources and improving IT competencies are essential to controlling these costs.

¹ Source: ESG and EMC

To solve this issue, EMC provides the following VNX storage efficiency technologies:

- Virtual Provisioning
- Fully Automated Storage Tiering Pools for Virtual Pools (FAST™ VP)
- Multicore FAST Cache
- File Deduplication and Compression
- Block Deduplication
- Block Compression

These features in their own way provide various ways to improve capacity usage, optimize performance, and align with green initiatives keeping simplicity in mind.

Audience

This white paper is intended for EMC customers, partners, and employees who are considering using the various EMC Storage Efficiency features. Minimal or no experience using EMC storage systems is expected.

Introduction

Storage efficiency is the ability to store and manage data while consuming the smallest amount of space with little to no performance impact, resulting in a lower total operational cost. Efficiency addresses the everyday demands of managing costs, reducing complexity, and limiting risk.

The challenge of every storage administrator is to maintain or improve service levels, performance, and capacity utilization, while reducing power consumption and Total Cost of Ownership (TCO).

This paper describes how organizations can reduce storage costs by increasing capacity utilization, simplifying storage management, reducing application downtime, and storing information in the most efficient way possible.

This white paper contains references to the following scenario: A storage administrator for a large financial company who is responsible for provisioning storage, providing the most efficient storage technologies, and supporting mission critical storage management applications, including architectural integrity and service levels agreements (SLAs).

Storage Efficiency Technologies overview

EMC offers multiple Storage Efficiency Technology solutions. Each feature complements the others, and provides a different approach to handling data and achieving the most efficient data services.

Unified Virtual Provisioning

Virtual Provisioning allows you to allocate storage using Thick and Thin LUNs within storage pools. Thick LUNs and Thin LUNs can reside within the same storage pool and share the pool's storage capacity. Both types of LUN can be easily provisioned, expanded, compressed and, deduplicated from the Unisphere interface.

Storage pool for block — Groups of available disk volumes that are used to allocate available storage to pool LUNs (Thick and Thin).

Storage pool for file — Groups of available disk volumes organized by Automatic Volume Management (AVM) that are used to allocate available storage to file systems. They can be created automatically with AVM, or manually by the user.

Thick LUN — A type of pool LUN where the allocated physical space is equal to the user capacity seen by the host server. Thick LUNs provide predictable high performance for applications primarily because Thick LUNs use direct LBA addressing instead of the fine-grained 8 KB mapping used by Thin LUNs. With Thick LUNs, all of the user capacity is reserved and allocated upon creation.

Thin LUN — A type of pool LUN where allocated physical space can be less than the user capacity seen by the host server. Thin LUNs can be sized to accommodate growth without regard for currently available assets, and provide on-demand storage to maximize storage utilization by allocating storage as needed.

File storage on VNX systems is a Thin-friendly environment. Thin-friendly is a term for file systems that do not pre-allocate all of the storage space during initialization. This term is also used for file systems that reuse deleted space before consuming additional storage. Both of these features improve capacity utilization in Thin provisioning.

Thin file system — A file system that lets you allocate storage based on long-term projections, while you consume only the File System resources that you currently need. NFS or CIFS clients and applications see the virtual maximum size of the File System of which only a portion is physically allocated.

FAST VP

Fully Automated Storage Tiering for Virtual Pools (FAST VP) is an advanced data service that automatically relocates data of pool-based LUNs at a sub-LUN level to optimal locations within a storage pool. FAST VP retains the most frequently-accessed data on fast, high-performance (more expensive) drives, and moves less frequently accessed data to lower-performance (less expensive) drives.

The available tiers when creating a storage pool are:

- Extreme Performance tier – Flash drives
- Performance tier – Serial Attach SCSI (SAS) drives
- Capacity tier – Near-Line SAS (NL-SAS) drives

FAST VP allows the assignment of different tiering policies within a pool. Tiering policies may be based on performance requirements, frequency of use, cost, and other considerations.

FAST VP uses an algorithm to make relocation decisions based on the activity level of each slice. The Unisphere interface allows users to set these LUN-level slice relocation policies with the following choices:

- **Highest Available Tier**
- **Auto-Tier**
- **Start High Then Auto-Tier (Recommended)**
- **Lowest Available Tier**
- **No Data Movement**

Multicore FAST Cache

Multicore FAST Cache is a large-capacity secondary cache that uses SAS Flash or SAS Flash 2 Flash drives positioned between the storage processor's DRAM-based primary cache and hard-disk drives (HDD).

Multicore FAST Cache extends the storage system's existing caching capacity for better system-wide performance. It achieves this by copying frequently accessed data to Flash drives, which are faster than HDDs, to boost system performance. Flash drives also provide a much larger, more scalable cache than the DRAM cache.

At a system level, Multicore FAST Cache makes the most efficient use of Flash drive capacity by using Flash drives for the most frequently accessed data in the storage system, instead of dedicating the drives to a particular application.

File Deduplication and Compression

VNX File Deduplication and Compression increases file storage efficiency by eliminating redundant data from the files stored in the file system, thereby saving storage space and money. For each file system, File Deduplication and Compression gives the Data Mover the ability to process files in order to compress them, and the

ability to identify and eliminate copies of files that are exact duplicates, reducing many copies to one and freeing up space. Deduplication functionality operates on whole files, and is applicable to files that are static or nearly static in a system-defined time.

Block Deduplication

Block Deduplication is the process of identifying duplicate data contained within a set of block or file storage objects and consolidating it so only one real copy of the data is used by many sources. This feature can provide significant space savings depending on the nature of the data. In the VNX2 storage systems, VNX Block Deduplication is enabled on a per pool LUN basis. These LUNs may be in the same pool as other non-deduplicated LUNs. Space freed by the deduplication process is returned to the pool for use by any type of LUN, deduplicated or not.

Block deduplication occurs at the 8 KB block level within the slices of the deduplication container in a pool.

VNX Block Deduplication for VNX File

In VNX OE for Block version 05.33.008.5.119 and VNX OE for File 8.1.8.119 VNX Block Deduplication can be used with VNX File. Enabling VNX Block Deduplication for a VNX File mapped storage pool or meta volume causes all underlying LUNs to become deduplication enabled thin LUNs. VNX Block Deduplication utilizes an 8 KB Fixed-Block Deduplication method of removing redundant data from all deduplication enabled LUNs within a Block Storage Pool. For more information on VNX Block Deduplication for File, management, and differences between VNX Block Deduplication for File and VNX File Deduplication and Compression refer to the *EMC VNX2 Deduplication and Compression* white paper on EMC Online Support.

Block Compression

Compression on VNX Block differs from VNX File Deduplication and Compression. Block Compression is a feature that lets you compress LUN data to free up storage space. Data compression analyzes the data on a disk and applies algorithms that reduce the size of repetitive sequences of bits inherent to some types of files.

Block Compression processes data in 64 KB increments, and compressed data is written to the LUN if at least 8 KB of the 64 KB can be saved. If the resulting savings from compression is less than 8 KB within the 64 KB chunk, the data is not compressed. The VNX does not compress data without sufficient savings.

Multi-dimensional storage efficiency

Storage efficiency is a multi-dimensional topic. Data reduction technologies like deduplication and compression, which reduce the data footprint, are fundamental components of storage efficiency. However, there are several additional technologies that come into play in this conversation. EMC is always looking to leverage the best breed of technologies in its products, whether they are developed internally or provided by a technology partner. Intel's Xeon E5 processor, used in the VNX2 series, is a perfect example of this. This processor uses its multi-core capabilities to scale performance across all cores and sockets. Storage efficiency technologies can be categorized along four dimensions. As illustrated in [Figure 2](#), when considering storage efficiency, is comprised of capacity optimization, performance optimization, simplicity, and green initiatives. Any given IT project at any stage in its lifecycle falls somewhere in this multi-dimensional diagram.



Figure 2. Multi-dimensional storage efficiency

It is important to remember how important each of these elements is to the end-user, keeping in mind that priorities may change over time. Coming back to the hypothetical scenario described in [Introduction](#), the financial company storage administrator plans to start the journey to the cloud. There different options to implement the cloud solution. One option is to start with a small number of virtual machines and plan for rapid growth. In this scenario the performance and simplicity of the infrastructure are most important aspects of the solution. Alternately, using a

large number of Tier 3 virtual machines shifts the focus to the capacity and simplicity of the storage required to support the machines, and ease of management.

Capacity optimization

Storage administrators must identify what is important to their organization if they hope to improve their storage environment.

Combining virtual server growth and data growth provides a fairly predictable result. An increasing amount of the data is accessed by virtualized applications. [Figure 3](#) shows how 71 percent of all storage capacity will be virtualized by 2016. The vast majority of this data is for applications that would exist regardless of whether data is running on a physical or virtualized server.




Figure 3. Worldwide external storage growth in virtualized environments²

With all the attention on server virtualization and predictions like the one shown above, it is critical to optimize the efficiency of the storage infrastructure that supports virtual machines.

Make sure to consider archive data in addition to active data, because it is very common to focus exclusively on active data. Storage administrators should certainly pay plenty of attention to active data, but ignoring archive data is ignoring the bulk of the data that must be managed.

Optimizing storage efficiency in capacity terms usually involves deploying a number of technologies in concert rather than any single solution.

² Source: IDC, Worldwide Storage and Virtualized x86 Environments External Storage Capacity by Hypervisor, 2011–2017 (TB)
IDC, Worldwide Storage and Virtualized x86 Environments 2011-2016 Forecast (Jul 2012)

When designing a cloud solution, consider the low infrastructure overhead such as power, cooling and space utilization. High capacity and high performance often counteract each other, so a well-balanced performance/capacity solution is important. Here are some key strategies to consider:

- Allocate capacity on demand using Thin Provisioning to minimize capacity and the associated power, cooling, and space requirements. Even with careful planning, it may be necessary to provision additional storage in the future. This may require application downtime depending on the operating systems involved. Thin Provisioning technology is designed to address these concerns. Thin LUNs can present more storage to an application than is physically available, eliminating the need for the time-consuming administrative work of deciding how to allocate drive capacity.
- Use compression to archive data to move it off of high-cost production resources.
- Use FAST VP to automatically relocate data from Thin LUNs to optimal locations within a storage pool.
- Deduplicate all active production data to reduce storage footprint and improve the efficiency of other performance features such as Multicore FAST Cache and FAST VP.

Capacity optimization is an umbrella term that refers to a number of different technologies, each of which has different cost benefit trades offs depending on the use case. Some of the features under this umbrella are file-level deduplication and block deduplication and compression.

Compression and deduplication are two different features with different functionalities. The question is: what technology or technologies are the best for a given type of data (file, block) at each point in the data's lifecycle?

The goal for the storage administrator is to choose the best technology to optimize capacity without impacting operational efficiency.

The information shown in [Table 1](#) explains how users must consider both the benefits and the costs associated with a given technology set for a given type of data at each point in its lifecycle. The results were taken from a 900GB EMC IT File System of traditional user data, such as office shares and files, home directories, and so on.

Table 1. Data deduplication and compression technologies

Technology	File deduplication	Block deduplication	Compression
Space savings	10%	20%	40-50%
Resource footprint	Low	High	Medium
File data	✓	✓	✓
Block data		✓	✓

The first row shows the typical space savings achieved by applying each of the technologies to a data set taken from EMC's own internal systems. The second row provides an estimate of the cost of applying that technology. This is not a dollar cost, but a runtime resources cost, which is determined by calculation of how much CPU and/or memory is required to realize the benefits in row two.

The checkmark symbols indicate the technology choices that are applicable to file and block data respectively. For file data, VNX applies an intelligent combination of File Deduplication and Compression because it provides the best return on investment for this type of data. For block data, VNX supports compression and block deduplication, but only one should be used at any given time.

The amount of capacity savings resulting from Block Compression depends on the type of data being stored. For example text files usually have large savings as compared to media files like .jpg files which usually already have native compression technologies. Unlike deduplication, compression does not depend on the presence of multiple copies of the same data. Deduplication occurs within one LUN, and among a group of LUNs. This is the main reason to deduplicate many LUNs together in the same pool.

These two features are designed for different use cases. Block Deduplication is designed for virtualized environments, which typically serve production and active data. Compression, on the other hand, is intended for archive data where the overhead due to decompressing data on subsequent accesses has less impact.

A standard virtual environment, as shown in [Figure 4](#), like the cloud storage solution discussed previously includes virtual desktops or virtual servers running Exchange, SQL and SharePoint. In this environment, there are multiple instances of the virtual machines with duplicate data, which makes it an ideal fit for deduplication.




Figure 4. Standard virtual environment

Block Deduplication is an ideal candidate for common virtual environments. The production environment will have substantial amounts of common blocks, such as multiple virtual desktop images, which are ideal targets for deduplication. Block Deduplication applies to *persistent desktops* and their internal home drives containing profiles and user data. Deduplication is also heavily used in multi-LUN redundant data settings. The cloud storage solution discussed in this paper will have substantial amounts of common blocks spanning from multiple LUNs, such as where there are test databases, or file systems with duplicate parts of files. At the same time multiple LUNs are provisioned from one original source. This environment has multiple LUNs which all derive from the same data set, thus sharing extensive amounts of common blocks. Block Deduplication is a key consideration for capacity optimization.

When enabled on a file system, VNX File deduplication works to remove duplicate files found within a common file system. This deduplication technique is referred to as Single-Instance Storage (SIS). SIS is not an EMC-specific deduplication technique, but rather is one of the many methods of data deduplication. By removing duplicate files, the amount of space needed for storage is reduced. The savings achieved while using File deduplication will vary as they depend on the similarity of the files being stored on the file system. Movie files, text files, documents, and pictures are good candidates for File Deduplication and Compression. For example, in home directory use cases where many users save the same picture or media files to their own directory, these files will be deduplicated together. At the same time, text files and documents achieve capacity savings due to compression.

Note: For more information on deduplication and compression, refer to the *EMC VNX2 Deduplication and Compression: VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, and VNX8000* white paper found on EMC Online support.

Performance optimization

Performance and capacity requirements for stored information change over time. In any given storage container, be it an array, a pool of drives within the array, a LUN, a file system, or a database, plotting the desired performance level of the data against the capacity provides a decay curve as shown with the multi-color line in Figure 5. Data typically moves down and occasionally up this curve over time. As data accumulates over time, it falls in activity rapidly, and after approximately three months the probability of hitting the same slice of data is only 10 percent.

Combine that with the fact that data accumulates over time, as illustrated by the exponential growth curve (green line) shown in Figure 5, and the net effect is that most growth is around capacity, not performance. This makes sense when considering how business applications generate and use data. The data that is in use today is accessed frequently, but as it ages, it is accessed less. Occasionally the activity associated with a piece of data may increase again at the end of the month or quarter, but ultimately it ends up on the tail of the curve where it needs to be stored because it contains valuable information but does not get accessed very often.

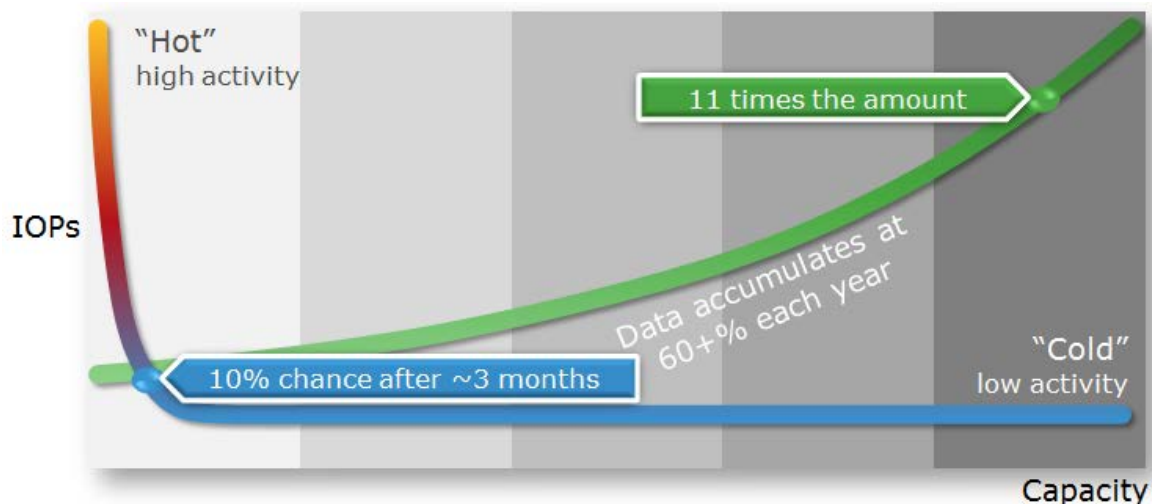


Figure 5. Typical data behavior

Virtual Provisioning, combined with the intelligence of FAST VP, provides a storage capability that matches the way that business applications use their storage.

Virtual Provisioning allows users to create LUNs and file systems that are supported by pools of storage containing any combination of Flash, SAS, and NL-SAS drives. With the FAST VP enabler installed, data requiring high performance can be given high performance, while the data that requires only capacity can be accessed efficiently and at lower cost. FAST VP allows the user to achieve this goal without over-spending on performance or capacity. The intelligence of FAST VP means that data is moved to the most appropriate type of storage as dictated by the tiering policy, and by the needs of the business applications accessing the data.

The VNX FAST Suite, which includes both FAST VP and Multicore FAST Cache, improves performance and maximizes storage efficiency by deploying EMC's Flash 1st

strategy. As a general rule, Multicore FAST Cache is used in cases where storage system performance needs to be improved immediately for burst-prone data. FAST VP is used to optimize TCO, moving data to the appropriate storage tier based on sustained data access and demands over time. Both features complement each other and play a crucial role in all heavy transactional workloads.

Multicore FAST Cache works with FAST VP, and is aware of the multi-tier data placement. Multicore FAST Cache is smart enough to ensure that storage system resources are not wasted by unnecessarily copying data to Multicore FAST Cache if it is already on a Flash drive. If FAST VP moves a slice of data to the Extreme Performance Tier, Multicore FAST Cache will not copy that slice of data; even it meets the promotion criteria. This ensures that storage system resources are not wasted copying data from one type of Flash drive to another.

If the user only has a limited number of Flash drives, EMC recommends using these drives for Multicore FAST Cache in the initial deployment phase. In most cases, Multicore FAST Cache, with a 64 KB granularity, offers better optimization of Flash in comparison with the 256 MB granularity offered by FAST VP.

Compounded efficiencies are attained when deduplication and compression are used in tandem with FAST VP and Multicore FAST Cache. When this is done, performance does not necessarily improve, but it makes the VNX performance features more efficient and more effective because of block sharing. Deduplication was expressly designed to reduce storage footprint, and reduce the TCO of FAST VP and Multicore FAST Cache while maintaining performance levels.

In terms of the interoperability between deduplication and FAST VP, the VNX2 series models allow these two features to work together to optimize tier utilization of the shared blocks. As activity increases on any of the virtual machines, the shared block is promoted to a faster tier. Because the block is shared, all machines utilize the same FAST VP policy across the shared data. This means that for all machines, the most active data is on the highest-performing tier, and the least active data is on the most cost-effective tier.

When adding Multicore FAST Cache into the equation, the VNX2 series ensures that frequently-accessed shared blocks maintain a high ranking for remaining in the Multicore FAST Cache. Because the blocks in the pool are shared, only a single FAST Cache entry per block is required, even though that region is accessed by multiple virtual machines. This is how deduplication increases the efficacy and efficiency of FAST Cache, each shared block from the pool results in a single entry in the FAST Cache.

Note: For more information on Multicore FAST Cache, FAST VP and Virtual Provisioning refer to the following whitepapers: *EMC VNX2 Multicore FAST™ Cache: VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, & VNX8000 – A Detailed Review*, *EMC VNX2 FAST™ VP: VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, & VNX8000 – A Detailed Review*, and *Virtual Provisioning for the VNX2 Series – VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, & VNX8000 – Applied Technology*, found on EMC Online Support.

Green Initiatives

It is also important to consider the implications of EMC storage efficiency technologies in terms of their environmental impact and cost savings.

Figure 6 compares the relative cost of the Extreme Performance tier (Flash drives), Capacity tier (SAS drives), and Capacity tier (NL-SAS drives). The two sets of columns on the left show the relative acquisition and running costs per gigabyte of these three storage technologies, by capacity. The two sets of columns on the right show the relative acquisition and running costs per I/O operation, by performance.

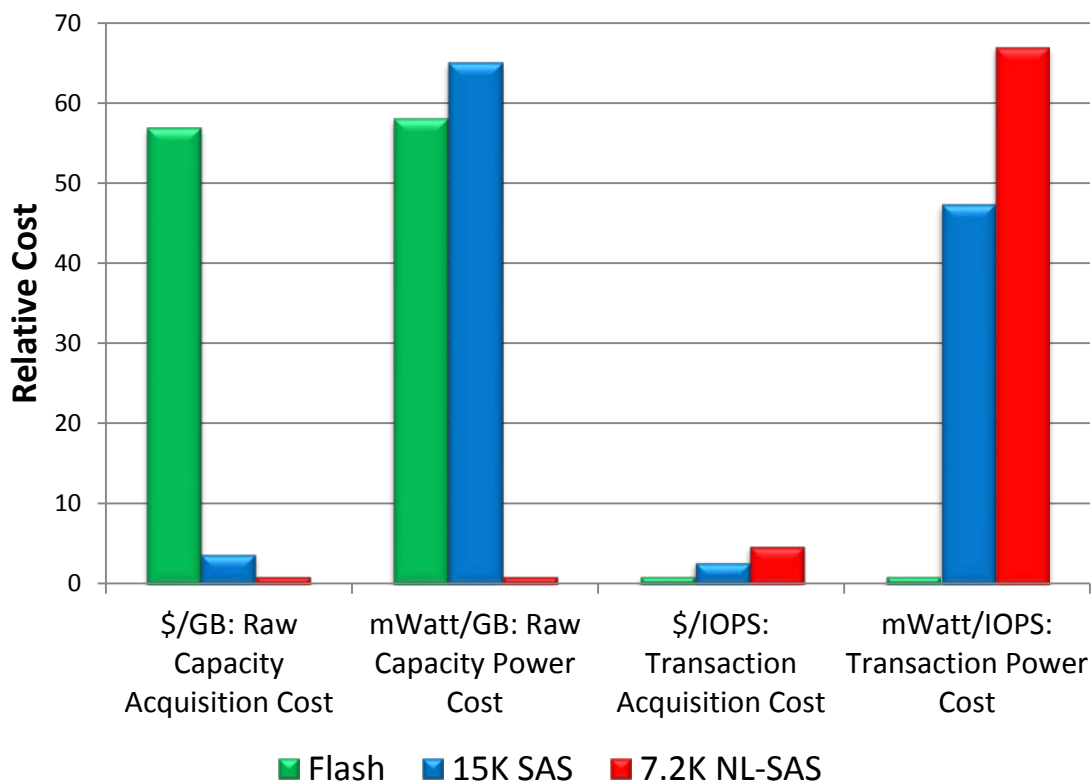


Figure 6. Disk drive efficiency considerations

If the driving metric is dollar per gigabyte, then high capacity NL-SAS are the best in terms of both acquisition and running costs. If transactional performance, that is IOPS, are the most important factor, then Flash drives are the best in terms of acquisition cost and environmental running costs.

Performance efficiency, as often measured by metrics such as dollars per I/O, is particularly important in transaction processing environments. The key consideration is the amount of hardware required to deliver a given number of IOPS at acceptable latency. Multi-core CPUs are a key technology in this space. The new Intel Xeon E5 processor makes the VNX2 faster and greener. This new 32 nm processor architecture provides multi-core capability that translates into significant scalability for each VNX model. The processor speeds range from 1.80 GHz to 2.7 GHz, and have four to eight

cores per socket. As a result, these processors provide the ability to accomplish more in a smaller footprint, consume less power, and keep systems cooler. This is the reason that the VNX2 series are faster than the previous generation of the EMC mid-tier platforms, and has sufficient processing power to change traditional storage management models with capabilities like FAST VP, and pool based storage provisioning. However, even better processors are not enough on their own. The processors work in combination with software architectures, like MCx™ in VNX, that are able to take maximum advantage of multiple cores to realize the full potential of fully automated, fully-optimized performance and capacity.

Note: For more information regarding the VNX2 systems series hardware and software architecture, refer to the *EMC VNX2 MCx™ Multicore Everything* white paper, and *Introduction to the EMC VNX2 Series: VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, & VNX8000 – A Detailed Review*, found on EMC Online Support.

Similarly, disk drives don't stand alone. They are packaged into arrays with controllers, buses, battery backup units, fans, power converters and so on. Therefore, while the power efficiency of the drives is important, the power efficiency of the array into which those drives are packaged is also a factor. A simple example is the fans in the array, which are speed-controlled based on temperature.

Finally, do not lose sight of the simple physical space requirements of the systems, and how much floor and rack space they consume. These are all important factors to consider, and have a positive impact on the environment.

In addition to the physical efficiency of a system, consider the logical efficiency of a virtual environment. One of the major benefits of virtual servers, such as the ones supported by VMware, is that they allow users to achieve greater efficiencies in the use of physical servers by consolidating multiple server instances to one physical host while preserving the individual identity of each client.

As shown in [Figure 7](#), one of the major benefits of an external storage array is the ability to consolidate multiple storage requirements into a high performance, shared storage infrastructure. VNX includes a feature called Virtual Data Movers (VDMs) that enable the presentation of multiple file servers from a single shared platform while maintaining security and identity between the individual file servers. This allows the consolidated file server platform to maintain the identity, authentication scope, and security boundaries of each individual file server.




Figure 7. Multi-tenancy for private cloud

Customers want more efficient IT solutions that cost less money to operate, consume less power, and help reduce carbon footprint. The VNX2 series models introduce new storage efficiency technologies, such as deduplication, and provide enhancements in Virtual Provisioning, FAST VP and FAST Cache. All these features deliver an energy-efficient design. For example, FAST VP delivers up to 40 percent more application performance at a 40percent lower cost, while requiring 87 percent fewer disk drives and 75 percent less power. Customers using Fast VP can reduce power consumption by more than 270 million KW per hour, enough to power 25,000 homes. In the upcoming decade, a highly efficient, well-designed green strategy for data management will be crucial to the success of any organization.

Simplicity

As computing environments grow, merge, evolve, and change, the risk of creating runaway complexity is very real. A more holistic approach to management, improved automation, and fewer, simpler user interfaces are key weapons in the fight against complexity. The VNX2 series models deliver new storage efficiency tools and methodologies that reduce complexity, and facilitate an easy-to-use experience.

Simplicity is how much effort it takes to manage a system day to day. Features such as the VNX Unisphere management interface make that as easy as possible.

Unisphere is completely web-enabled for remote management of the storage environment. With customizable view blocks and sortable tables, critical information is never more than a few clicks away. VNX Unisphere provides a flexible, integrated experience for managing the VNX2 series storage systems. Unisphere's unprecedented ease of use is reflected in its intuitive task-based controls, customizable dashboards, context-based application management, and single-click access to real-time support tools and online customer communities.

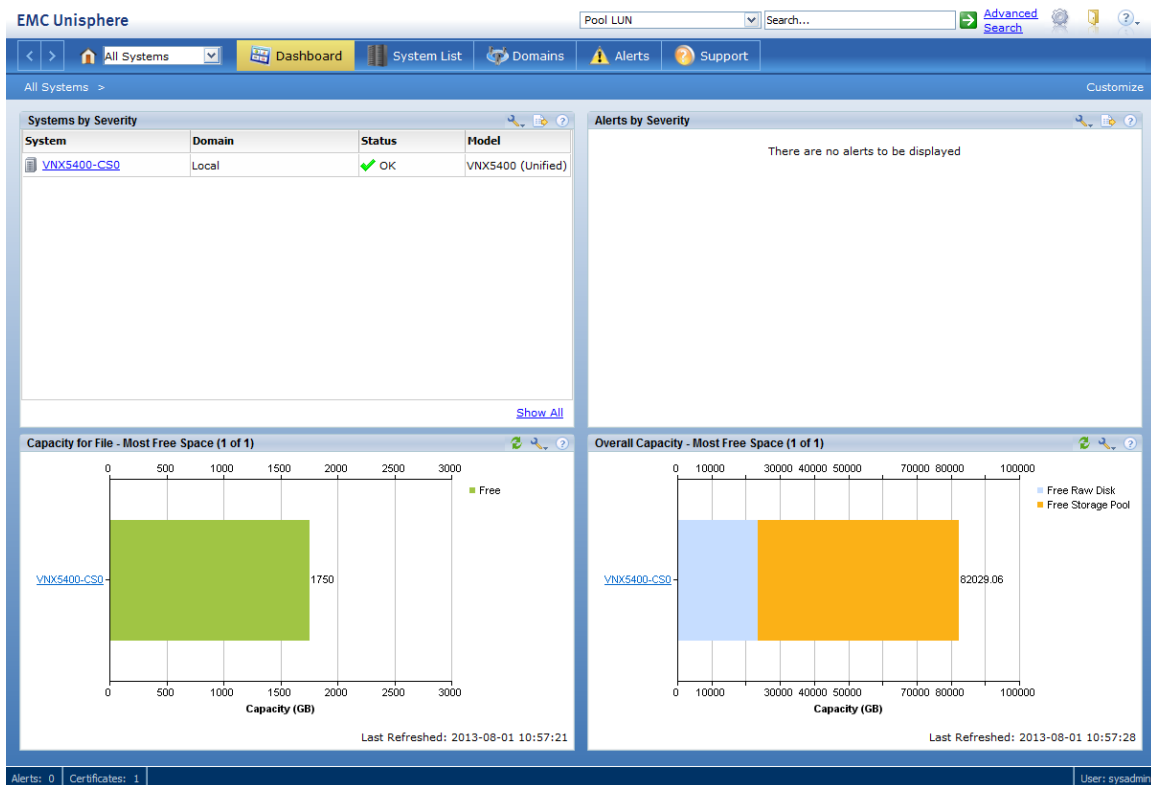


Figure 8. VNX Unisphere's Enterprise Dashboard

Unisphere is a single management interface for all unified systems. It allows you to manage all the storage efficiency technologies described in this paper. For example, users can provision storage LUNs (Thick or Thin) with the LUN Provisioning Wizard, schedule FAST VP relocations, change tiering policies, or enable, configure, and disable deduplication and compression. Multicore FAST Cache can be created easily. Unisphere Analyzer provides the ability to monitor performance metrics.

VNX Unisphere integrates with other EMC products such as Atmos, Isilon, RecoverPoint/SE, Data Protection Advisor, and more.

From a strategic perspective, simplicity is flexibility that enables agility without cost. This enables storage administrators to deliver the services their business needs when required. It is about having choice, now and in the future as business needs evolve. A flexible storage infrastructure can make it easier to decide what type and how much storage infrastructure to deploy, as well as when and how storage is consumed.

Note: For more information regarding the VNX Unisphere management interface, refer to the *EMC Unisphere: Unified Storage Management Solution for the VNX2 Series: VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, & VNX8000 – A Detailed Review* white paper, found on EMC Online Support.

Conclusion

Organizations of every type rely on the timely retrieval of information to facilitate transactions and decision making. The VNX2 series models are equipped with storage efficiency technologies that help simplify the many IT challenges that companies are facing. Typical issues facing organizations include annual double-digit data growth, ever increasing IT budgets, and over/under staffing. The VNX series is designed to maximize capacity optimization and performance optimization, while promoting green initiatives and ensuring simplicity of use. The VNX series is expressly configured to take advantage of the latest innovations in Flash drive technology, maximizing the storage system performance and efficiency, while minimizing cost per gigabyte.

References

- *Introduction to the EMC VNX2 Series: VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, & VNX8000 – A Detailed Review*
- *EMC Unisphere: Unified Storage Management Solution for the VNX2 Series: VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, & VNX8000 – A Detailed Review*
- *EMC VNX2 Deduplication and Compression: VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, & VNX8000*
- *EMC VNX2 FAST™ VP: VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, & VNX8000 – A Detailed Review*
- *EMC VNX2 MCx™ Multicore Everything*
- *EMC VNX2 Multicore FAST™ Cache: VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, & VNX8000 – A Detailed Review*
- *Virtual Provisioning for the VNX2 Series: VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, & VNX8000 – Applied Technology*